

CERTIFICATE OF TRANSMISSION

I hereby certify that this correspondence is being transmitted by either electronic submission using the EFS WEB submission system, fax to the U.S. Patent and Trademark office to fax number 571-273-8300, or is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450 to on May __19__, 2006.

_____/David J McKenzie/
Attorney for Applicant

PATENT

Docket No. DE920000124US1

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appellant:	Christian Bolik, et al.)
)
Serial No.:	10/015,825)
) Group Art
Filed:	December 10, 2001) Unit: 2145
)
For:	Method and System for Scalable, High)
	Performance Hierarchical Storage)
)
Examiner:	Ajay M. Bhatia)
)

APPELLANT'S AMENDED APPEAL BRIEF

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

On January 19, 2005, Appellant filed a timely Notice of Appeal from the Final Office Action mailed September 20, 2005. Along with the Notice of Appeal, Appellant filed a Petition for Extension of Time and a fee for a one month extension. Appellant appeals from the rejection of all pending claims.

This Brief is being filed under the provisions of 37 C.F.R. § 41.37. This Brief also responds to the Notification of Non-Compliant Appeal Brief mailed 4/19/2006. Specifically, this Brief responds the Examiner's arguments regarding the summary of claimed subject matter by specifically referencing the independent claims and means plus function claim in the summary and throughout the Brief and by adding a summary of the independent claims. The filing fee set forth in 37 C.F.R. § 41.20(b)(2) of Five Hundred Dollars (\$500.00) was previously being submitted. The Commissioner is hereby authorized to charge payment of any additional fees associated with this communication, or to credit any overpayment, to Deposit Account No. 09-0466.

1. REAL PARTY IN INTEREST

The real party in interest is the assignee, International Business Machines Corporation, Armonk, New York.

2. RELATED APPEALS AND INTERFERENCES

There are no related appeals and interferences.

3. STATUS OF CLAIMS

In the Office Action mailed September 20, 2005 ("Final Office Action"), Claims 1-20 were rejected. The Final Office Action rejected claims 1-20 under 35 U.S.C. § 102(e) as being anticipated by U.S. Patent No. 6,269,382 to Cabrera, et. al. (hereinafter "Cabrera"). The Appellants responded by submitting a response to the Final Office Action on November 21, 2005

(hereinafter “Final Office Action Response” or “FOAR”) which amended Claims 1, 15, 19, and 20 and cancelled Claim 2. The Advisory Action mailed December 21, 2006 (hereinafter “Advisory Action”) stated that the “Applicant’s reply has overcome the following rejection(s): 35 U.S.C. 102(e) Cabrera claims 1, 3-20.” Advisory Action, p.2. However, the Advisory Action also alleged that the “Applicant has submitted the same argument that were (sic) addressed in the final rejection” and referenced the rejection made in the Final Office Action. *Id.* at p. 3. Notwithstanding the statement in the Advisory Action that the Applicants have overcome the rejections to claims 1, 3-20, presumably claims 1, 3-20 stand rejected under 35 U.S.C. § 102(e).

Appellant respectfully appeals the rejection of Claims 1, 3-20.

4. STATUS OF AMENDMENTS

Appellant amended Claims 1, 15, 19, and 20 and cancelled Claim 2 in the Final Office Action Response and the amendments to the claims were entered effective December 21, 2005.

5. SUMMARY OF CLAIMED SUBJECT MATTER

The claimed subject matter deals with Hierarchical Storage Management (“HSM”). By way of background, the present invention comprises a method and system for a scalable high performance HSM. Specification at p. 1, ll. 6-8. A typical HSM frees up more expensive local storage, such as magnetic disk storage, that is limited by migrating data files to less expensive remotes storage, such as tape storage, optical storage, etc. that is virtually infinite. *Id.* at p. 1, ll. 14-17. Migration involves copying a file on the managed system (local storage) to remote storage and replacing the file on the local storage with a small “stub” file that is indistinguishable

to the user from the actual data file. *Id.* at p. 1, ll. 14-21. A typical HSM finds candidate data files for migration by traversing the complete file structure of the local storage to identify the candidate files before migrating the candidate data files. *Id.* at p. 4, ll. 5-6. Selection of candidate files for migration is based on a migration policy with parameters such as file size, length of time a file has been stored, whether a file has been previously migrated, etc. *Id.* at p. 2, ll. 6-12. Candidate files are migrated to the remote storage at a particular time or in response to the local storage reaching or nearing capacity. *Id.* at p. 2, ll. 14-19.

Complete traversal of the file system impedes scalability because resource requirements and duration of the traversal grow logarithmically with the number of files in a file system. *Id.* at p. 4, ll. 8-10, p. 9, ll. 9-10, Fig. 2. The claimed invention overcomes the problems associated with the complete traversal of the file system by scanning the local storage to find migration candidate files until a prespecified number of candidate files are found. FOAR at claim 1; Spec. at p. 9, ll. 11-16, p. 5, ll. 8-10. Once the number of identified candidate files reaches the prespecified number, scanning may be terminated or temporarily halted. Spec. at p. 13, ll. 1-4. The migration process may use the partial list of candidate files from a scanning that was halted after reaching the prespecified number of candidate files and need not wait until the entire file structure is traversed, thus making the candidate determination process practically independent of the number of files on a file system. *Id.* at p. 13, ll. 13-16. The claimed invention works equally well for files in a 'migrated' state and for files in a 'pre-migrated' state where data files are stored on the HSM server and an identical copy is stored on the local storage. *Id.* at p. 7, ll. 4-7.

Claim 1 presents a method for managing an HSM environment. FOAR at claim 1; spec. at p. 9, ll. 11-12. The HSM environment includes at least one HSM server and at least one file

server having stored a managed file system. FOAR at claim 1; spec. at p. 9, ll. 2-8, Fig. 1. The at least one HSM server and the at least one file server are interconnected via a network and digital data files are migrated temporarily from the at least one file server to the at least one HSM server. *Id.* The method includes the steps of providing at least one list for identifying candidate data files to be migrated and prespecifying a scanning scope determined by a number of candidate data files. FOAR at claim 1; spec. at p. 9, ll. 12-14, Fig. 3. The method includes scanning the managed file system until having reached the prespecified number of migration candidate data files. FOAR at claim 1; spec. at p. 9, ll. 14-16, Fig. 3. The method includes selecting migration candidate data files according to at least one attribute and recording the selected migration candidate data files in the provided at least one list for identifying candidate data files. FOAR at claim 1; spec. at p. 9, ll. 17-21, Fig. 3. The method also includes migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server. FOAR at claim 1; spec. at p. 10, ll. 4-8, Fig. 3.

Additionally, an embodiment of the present invention includes an apparatus claimed in means plus function format under 35 U.S.C. § 112, sixth paragraph. Examples of the structure, material, or acts corresponding to the means of Claim 15 are referenced below. Claim 15 presents an HSM system including at least one HSM server and at least one file server having stored a managed file system. *Id.* at claim 15. The at least one HSM server and the at least one file server are interconnected via a network, where data files are migrated temporarily from the at least one file server to the at least one HSM. *Id.* The HSM system includes means for prespecifying a scanning scope determined by a number of candidate data files. *Id.* An HSM

server is an example of means for prespecifying a scanning scope. Spec. at p. 9, ll. 5-8, Fig. 1. Another example of means for prespecifying a scanning scope is the prespecifying scanning scope step 200 of the flowchart shown in Figure 3. *Id.* at p. 9, ll. 11-14.

The HSM system includes means for scanning the file system until having reached the prespecified number of migration candidate files and for identifying candidate data files to be migrated. FOAR at claim 15. An HSM server is an example of means for scanning the file system until having reached the prespecified number of migration candidate files and also for identifying candidate data files to be migrated. Spec. at p. 9, ll. 5-8, Fig. 1. The serial scanning step 201 and the determining file candidate step 202 of the flow chart shown in Figure 3 are also examples of means for scanning the file system until having reached the prespecified number of migration candidate files and for identifying candidate data files to be migrated. *Id.* at p. 9, ll. 14-19, Fig. 3.

The HSM system includes means for monitoring the managed file system. FOAR at claim 15. The HSM system is an example of means for monitoring the managed file system. Spec. at p. 9, ll. 5-8, Fig. 1. The monitoring the file status step 205 and the determining the current status of the file system step 206 of the flowchart shown in Figure 3 are also examples of means for monitoring the managed file system. *Id.* at p. 9, ll. 25-26, Fig. 3.

The HMS system includes means for migrating candidate data files to the HSM server. FOAR at claim 15. The HSM system is an example of means for migrating candidate data files to the HSM server. Spec. at p. 9, ll. 5-8, Fig. 1. In addition, the initiating automigration step 207 and the performing automigration and assigning a unique identifier step 208 of the flowchart shown in Figure 3. *Id.* at p. 9, l. 26 to p. 10, l. 8, Fig. 3.

The HSM system includes means for reconciling the managed file system. FOAR at claim 15. The HSM system is an example of means for reconciling the managed file system. Spec. at p. 9, ll. 5-8, Fig. 1. The reconciliation steps 301, 302, and 303 of the flowchart shown in Figure 4 are further examples of reconciling the managed file system. *Id.* at p. 10, ll. 9-17.

Claim 19 presents a data processing program for execution in a data processing system that includes software code portions for performing a method. FOAR at claim 19; spec. at p. 9, ll. 11-12. The method includes the steps of providing at least one list for identifying candidate data files to be migrated, prespecifying a scanning scope determined by the number of candidate data files, and scanning the managed file system until having reached the prespecified number of migrated candidate files. FOAR at claim 19; spec. at p. 9, ll. 12-16, Fig. 3. The method also includes selecting migration candidate data files according to at least one attribute, recording the selected migration candidate data files in the provided at least one list for identifying candidate data files, and migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server. FOAR at claim 19; spec. at p. 9, ll. 17-21, p. 9, l. 26 to p. 10, l. 8, Fig. 3.

Claim 20 presents an article of manufacture that includes a program storage medium readable by a processor and embodying one or more instructions executable by the processors to perform a method. FOAR at claim 20; spec. at p. 9, ll. 11-12. The method includes providing at least one list for identifying candidate data files to be migrated, prespecifying a scanning scope determined by a number of candidate data files, and scanning the managed file system until having reached the prespecified number of migration candidate data files. FOAR at claim 20; spec. at p. 9, ll. 12-21, Fig. 3. The method includes selecting migration candidate data files

according to at least one attribute, recording the selected migration candidate data files in the provided at least one list for identifying candidate data files, and migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server. FOAR at claim 19; spec. at p. 9, l. 26 to p. 10, l. 8, Fig. 3.

6. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

The following issues are presented for review:

- I. Did the Examiner fail to establish a *prima facie* case of anticipation under 35 U.S.C. § 102(e) for claims 1, 3-20 where the limitations of the claims are not taught or suggested within the cited reference?
- II. Did the Examiner improperly reject Claim 3 because Cabrera does not disclose scanning until having reached a prespecified amount of data?
- III. Did the Examiner improperly reject Claim 4 because Cabrera does not disclose resuming scanning at a location of the managed file system where a previous scanning left off?
- IV. Did the Examiner fail to establish a *prima facie* case of anticipation under 35 U.S.C. § 102(e) for claims 1, 3-20 where the Examiner failed to give proper weight to the limitation added by amendment further clarifying that scanning of a file system is limited to a prespecified number?

7. ARGUMENT

I. CLAIMS 1 AND 3-20 ARE NOT ANTICIPATED BY CABRERA UNDER 35 U.S.C. § 102(E)
BECAUSE ALL OF THE LIMITATIONS OF THE CLAIMS ARE NOT TAUGHT OR SUGGESTED
BY CABRERA.

The Prior Art. Cabrera discloses pre-migrating files that normally would not be migrated under a traditional migration policy. Cabrera at Abstract, col. 10, l. 30 to col. 11, l. 6. Cabrera discloses ‘pre-migration’ of files by copying a file to remote file storage but maintaining the source file in place in the local storage rather than replacing it with a stub as in the traditional migration process. *Id.* at col. 11, ll. 25-46. Cabrera also discloses recording which files are pre-migrated and then deleting the file on the local storage and replacing it with a stub file or truncating the file when the time for migration of the file is reached. *Id.* at col. 11, l. 47 to col. 12, l. 6, col. 5, ll. 1-26. Cabrera discloses scanning for candidate files for pre-migration by selecting files that would meet the migration policy in every way except the age of the file on the local storage. *Id.* at col. 10, ll. 30-44. Cabrera does not teach a method of limiting the scanning scope to less than scanning the entire local storage tree to find candidate files.

The Examiner does not make out a *prima facie* case of anticipation under 35 U.S.C. § 102(e) because Cabrera does not teach prespecifying a scanning scope determined by a number of candidate files required in claim 1. “A claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference.” *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631 (Fed. Cir. 1987). For a *prima facie* case of anticipation, each and every element of the claimed invention must be identically disclosed in a single prior art reference; and those elements must be arranged or connected together in a single reference in the same way as specified in the patent claim.

Lindemenn Maschinenfabrik GmbH vs. American Hoist and Derick Co., 730 F2d 1452, 221 USPQ 481, 485 (Fed. Cir. 1984). The Examiner has not met the burden of providing *prima facie* evidence of anticipation by demonstrating that each and every element of the claims is identically disclosed in Cabrera.

A. Cabrera Does Not Teach Prespecifying a Scanning Scope Determined by a Number of Candidate Data Files as Required by Claim 1.

Amended Claim 1 states:

- I. A method of managing a hierarchical storage management (HSM) environment, the environment including at least one HSM server and at least one file server having stored a managed file system, wherein the at least one HSM server and the at least one file server are interconnected via a network and wherein digital data files are migrated temporarily from the at least one file server to the at least one HSM server, the method comprising:
 - providing at least one list for identifying candidate data files to be migrated;
 - prespecifying a scanning scope determined by a number of candidate data files;
 - scanning the managed file system until ~~the scanning scope is reached~~ having reached the prespecified number of migration candidate data files;
 - selecting migration candidate data files according to at least one attribute;
 - recording the selected migration candidate data files in the provided at least one list for identifying candidate data files; and
 - migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server.

In an Office Action mailed January 26, 2005 (hereinafter “Office Action 1”), the Examiner argues that Claim 1 is anticipated by Cabrera. Office Action 1 at pp. 2-3. The Examiner mistakes pre-migrating files based on a pre-migration policy with a prespecified scanning scope that limits the number of files on a local storage that are scanned to find candidate files that meet a migration policy. Cabrera discloses selecting candidate files for pre-migration based on a pre-migration policy and then pre-migrating the candidate files by copying the file to remote file storage while leaving a copy of the file intact on the local storage. Cabrera at col. 10, l. 30 to col. 12, l. 6. Claim 1 discloses scanning only a portion of the files of a local storage to select candidate files for migration, where the number of files scanned is prespecified in a scanning scope, and then migrating the selected candidate files. Spec. at Claim 1, p. 9, ll. 11-21, p. 10, ll. 4-8.

The Examiner cites column 9, line 53 to column 10, line 17, column 4, lines 8-19, and column 10, lines 45-53 as evidence of anticipation. *Id.* Cabrera at column 9, line 53 to column 10, line 17 merely describes a traditional migration process of identifying migration candidates, migrating data, recording the migration, and freeing local storage. Cabrera at col. 9, l. 53 to col. 10, l. 17. Cabrera states only that the process “can use any number of mechanisms to identify which file should be migrated,” *id.* at col. 9, ll. 60-62, but does not indicate any method for scanning the local storage to find the candidate files. Cabrera at column 4, lines 8-19 merely describes how an HSM handles data requests for data that has been migrated by using an additional ‘attribute’ added to the migrated or pre-migrated files and allowing the HSM to handle data requests. *Id.* at col. 4, l. 8-19. The cited section does not discuss scanning the local storage for candidate files. *Id.*

Cabrera at column 10, lines 45-53 lists typical parameters for a migration policy including minimum candidate file size, age that a candidate file has been stored, a minimum amount of free space desired on the local storage, and files exempt from migration. *Id.* at col. 10, ll. 45-53. None of the cited text from Cabrera mentions any method of scanning and does not describe prespecifying a scanning scope. Instead, Cabrera only discusses identifying candidate files for pre-migration using a migration policy different than the policy used for full migration. *Id.* at col. 10, ll. 30-44. The pre-migration policy allows identification of candidate files where the files have not met the time criteria of the full migration policy. *Id.* Cabrera does not describe a scanning process to find pre-migration candidate files but implies only some sort of scanning process to identify the files. *See generally id.*

After the Applicants pointed out that Cabrera does not disclose prespecifying a scanning scope in an Office Action Response dated June 27, 2006 (hereinafter “Office Action Response 1”), the Examiner argued in the Final Office Action that Cabrera “as applied to the limitation as the scanning scope as all files that are in this pre-migrated state, and the prespecifying scope is the criteria for files to be designated as a pre-migration file.” Final Office Action at p. 2. The Examiner points to Cabrera at column 10, lines 25-30, column 10, lines 46-53, column 4, lines 51-55, column 5, lines 1-25, and column 5, lines 36-50 as evidence. *Id.*

Again, the Examiner confuses a scanning scope with migration policy. Cabrera at column 10, lines 25-30 merely states that the Cabrera invention identifies candidate files for pre-migration. Cabrera at col. 10, ll. 25-30. The cited section mentions nothing as to how the Cabrera invention scans the local storage. *See id.* Cabrera at column 10, lines 46-53 merely lists parameters for a traditional migration policy. *Id.* at col. 10, ll. 46-53. Again, the cited text

mentions nothing about the scanning process, but only states a typical migration policy. *See id.* Cabrera at column 4, lines 51-55 states that once candidate files for pre-migration are identified, the files are ‘pre-migrated’ to the remote storage. *Id.* at col. 4, lines 51-55. Again, the cited text does not say how the scanning for candidate files is accomplished. *See id.* Cabrera at column 5, lines 1-25 states that when a pre-migrated file meets the criteria for full migration, the file on the local storage is truncated, leaving a file that is useful as a traditional stub file. *Id.* at col. 5, ll. 1-25. The cited text only describes the advantages of this truncation process and only states that the HSM “waits until the file actually meets the migration criteria, . . .” *Id.* at col. 5, ll. 1-3. There is no mention of how the HSM scans the local storage to determine if the migration policy has been met by the files, but only an inference that the HSM somehow monitors the age of the pre-migrated files. *See id.* Cabrera at column 5, lines 36-50 discusses the situation where a pre-migrated file is accessed prior to the time when the file complies with the full migration policy. *Id.* at col. 5, ll. 36-50. The cited text states that when a pre-migrated file is accessed, the pre-migration is invalidated leaving the original file intact and the HSM deletes the copy on the remote storage. *Id.* The cited text has nothing to do with scanning the local storage.

The cited text in Cabrera does not deal with scanning scope but only creates a migration policy for pre-migration that is broader than a migration policy for full migration. Some type of scanning is only implied in the cited text based on statements that pre-migration and migration candidate files are identified. The cited text does not disclose how the HSM scans the local storage to find the candidate files. The argument that pre-migrating some files before full migration is not in any way equivalent to limiting the scope of scanning of the local storage for candidate files by prespecifying a scanning scope. The Applicants respectfully assert that

Cabrera does not teach every limitation of Claim 1 because Cabrera does not teach limiting scanning for candidate files by prespecifying a scanning scope determined by a number of candidate data files. The Applicants assert that Claim 1 is in condition for allowance and request that the rejection of Claim 1 by the Examiner be reversed.

B. Cabrera Does Not Disclose Scanning the Managed File System Until Having Reached the Prespecified Number of Migration Candidate Data Files as Required by Claim 1.

The Examiner is incorrect in asserting that Cabrera discloses scanning the managed file system until having reached the prespecified number of migration candidate data files, as required by Claim 1. The Examiner asserts that scanning the managed file system until the scanning scope is reached is “the files that are currently pre-migrated which are scanned through to find files which meet the migration requirements so that they can be removed from the managed file system.” Final Office Action at pp. 2-3. Scanning for candidate files for pre-migration requires scanning the local storage (managed file system) using a migration policy specific to the pre-migration process. Cabrera at col. 10, ll. 24-44. Cabrera does not disclose any scanning process, but only discloses that candidate files are identified. Cabrera then discloses that the pre-migrated files are fully migrated by truncating the corresponding file on the local storage but does not disclose any process of scanning the local storage to identify the pre-migrated files available for full migration. *Id.* at col. 12, ll. 22-68.

Pre-migrating files is a migration process that simply uses a different migration policy than full migration. A migration policy is a set of criteria for selecting a candidate file. Scanning uses the selection criteria of a migration policy to traverse the directories of the local storage and

identify files meeting the criteria. Scanning is not equivalent to pre-migration. Cabrera does not disclose scanning the managed file system until having reached the prespecified number of migration candidate data files. The Applicants assert that the Examiner has not established a *prima facie* case of anticipation by Cabrera because Cabrera does not teach every element of Claim 1. The Applicants assert that Claim 1 is in condition for allowance and that the Examiner's rejection is improper and should be reversed.

II. CLAIM 3 IS ALLOWABLE BECAUSE CABRERA DOES NOT DISCLOSE SCANNING UNTIL HAVING REACHED A PRESPECIFIED AMOUNT OF DATA.

Claim 3 is allowable because the Examiner is incorrect in asserting that Cabrera discloses scanning the local storage until having reached a prespecified amount of data. The Examiner asserts that Cabrera discloses scanning for candidate files until reaching a prespecified amount of data and cites Cabrera at column 10, line 65 to column 11, line 6. Office Action 1 at p. 3. The cited text in Cabrera states: “[f]or example, the present invention enables the system to examine the size of a file, and if that size is smaller than the size specified by the system manager, pre-migrate the file from the local storage to remote storage if the migration of a smaller-than-specified file could be justified. Such justification might include the introduction of additional storage space into the local storage. This migration violates the exemplary parameter 1 [minimum file size to be eligible for migration] of the migration policy above.” Cabrera at col. 10, l. 65 to col. 11, l. 6. The cited text only discloses examination of the size of a particular file and indicates that a file that would not meet the full migration policy may be pre-migrated. This has nothing to do with scanning a local file system and identifying candidate files until the amount of data comprising the identified candidate files reaches a prespecified amount of data, as

required by Claim 3. The Applicants respectfully assert that the Examiner has not established a *prima facie* case that Claim 3 is anticipated by Cabrera. The Applicants assert that Claim 3 is in condition for allowance and that the Examiner's rejection is improper and should be reversed.

III. CLAIM 4 IS ALLOWABLE BECAUSE CABRERA DOES NOT DISCLOSE RESUMING SCANNING AT A LOCATION OF THE MANAGED FILE SYSTEM WHERE A PREVIOUS SCANNING LEFT OFF.

Claim 4 is allowable because the Examiner incorrectly asserts that Cabrera discloses resuming scanning of the managed file system (local storage) where a previous scanning left off. The Examiner cites Cabrera at column 10, lines 45-53 to assert that Cabrera discloses resumption of scanning where scanning previously left off. Office Action 1 at p. 3. The cited text in Cabrera merely describes parameters of a typical migration policy of a traditional migration system. Cabrera at col. 10, ll. 45-53. The parameters include 1) a minimum file size for migration, 2) age of the file on the system, 3) a desired percentage of free space on the local storage, and 4) files exempt from migration. *Id.* The cited text only describes the criteria used during a scanning process to identify candidate files. The cited text does disclose a resumption of scanning for candidate files from a location where scanning previously left off. The Applicants respectfully assert that the Examiner has not established a *prima facie* case that Claim 4 is anticipated by Cabrera. The Applicants assert that Claim 4 is in condition for allowance and that the Examiner's rejection is improper and should be reversed.

IV. CLAIMS 1, 3-20 SHOULD BE ALLOWED BECAUSE THE EXAMINER ADMITTED THAT THE APPLICANTS HAVE OVERCOME THE REJECTION TO CLAIMS 1, 3-20 AND THE EXAMINER DID NOT GIVE ANY WEIGHT TO THE AMENDMENTS TO THE CLAIMS.

Claims 1, 3-20 should be allowed because the Examiner admitted in the Advisory Action that the Applicants had overcome the Examiner's rejections of the claims and because the Examiner improperly rejected the claims without considering the arguments of the Final Office Action Response and amendments to the claims. In the Advisory Action, under the heading of "Amendments" the Examiner indicated that the "Applicant's reply has overcome the following rejection(s): 35 U.S.C. 102(e) Cabrera claims 1, 3-20." Advisory Action at p. 2. The Advisory Action states that the amendments to the claims have been entered. *Id.* However, the Advisory Action adds that the request for reconsideration has been considered but has not placed the claims in condition for allowance and provides only a terse statement that the Applicant made the same arguments as were previously made. *Id.* at pp. 2-3.

The amendments to the claims clarify that the claimed invention is directed toward scanning of the local storage to find migration candidate files and thus further distinguish Cabrera. Final Office Action Response at p. 10-11. The Examiner's reply in the Advisory Action gives no weight to the amendments and additional arguments made in the Final Office Action.

Dependent Claims 5-12 depend from Claim 1 and because Claim 1 is in condition for allowance, the Applicants assert that Claims 5-12 are allowable. In addition, Claims 13-20 are similar in scope to Claims 1-12 and the arguments made above apply equally as well to Claims 13-20 and the Applicants assert that Claims 13-20 are allowable. In view of the foregoing, the Examiner has not properly established *prima facie* anticipation by Cabrera of claims 1, 3-20. Appellant respectfully requests reversal of the Section 102 rejection and allowance of claims 1, 3-20. Appellant submits that the foregoing arguments further establish novelty of the present

invention. Reversal of the rejections and allowance of the pending claims is respectfully requested.

SUMMARY

In view of the foregoing, each of the claims on appeal has been improperly rejected. Reversal of the Examiner's rejection and allowance of the pending claims 1, 3-20 is respectfully requested.

Respectfully submitted,

_____/David J McKenzie/_____

David J. McKenzie
Reg. No. 46,919
Attorney for Appellant

Date: May 19, 2006
8 East Broadway, Suite 600
Salt Lake City, UT 84111
Telephone (801) 994-4646
Fax (801) 322-1054

8. Claims Appendix

Claims involved in the appeal

1. A method of managing a hierarchical storage management (HSM) environment, the environment including at least one HSM server and at least one file server having stored a managed file system, wherein the at least one HSM server and the at least one file server are interconnected via a network and wherein digital data files are migrated temporarily from the at least one file server to the at least one HSM server, the method comprising:
 - providing at least one list for identifying candidate data files to be migrated;
 - prespecifying a scanning scope determined by a number of candidate data files;
 - scanning the managed file system until having reached the prespecified number of migration candidate data files;
 - selecting migration candidate data files according to at least one attribute;
 - recording the selected migration candidate data files in the provided at least one list for identifying candidate data files; and
 - migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server.
2. Canceled
3. The method according to claim 1, wherein the scanning scope is determined by the total amount of data for the candidate data files and wherein the managed file system is

scanned until having the prespecified amount of data.

4. The method according to claim 1, wherein the scanning of the managed file system is resumed at a location of the managed file system where a previous scanning is left off, and continued accordingly.
5. The method according to claim 1, wherein replacing a migrated data file in the managed file system by a stub file providing at least information about the location of the migrated data file on the HSM server.
6. The method according to claim 1, further comprising monitoring a current state of the managed file system and initiating automigration dependent on the monitored current state of the managed file system.
7. The method according to claim 6, comprising the further steps of automigrating candidate data files with respect to the list for identifying candidate data files and assigning a unique identifier to each of the migrated candidate data files.
8. The method according to claim 7, wherein the unique identifier is specific to the underlying file system allowing direct access to a migrated data file.

9. The method according to any of claim 6, wherein providing two lists for identifying candidate data files, whereby the first list is generated and/or updated by a scanning process and whereby the second list is used by a automigration process, and whereby the automigration process gathers the first list from the scanning process when all candidate data files of the second list are migrated.
10. The method according to any of claim 9, wherein the automigration process is performed by a master/slave concept where the master controls the automigration process and selects at least one slave to migrate candidate data files provided by the master.
11. The method according to claim 1, comprising the additional steps of ranking and sorting the candidate data files contained in the at least one list for identifying candidate data files, in particular with respect to the a file size and/or time stamp of the data files contained in the at least one list for identifying candidate data files.
12. The method according to claim 1, wherein the scanning of the managed file system is initiated dependent on expiration of a prespecified wait interval or initiated by the automigration process.

13. A method of reconciling a managed file system migrated from a file server to an hierarchical storage management (HSM) server via a network in accordance with the method according to any of claims 7 to 12, with a current state of the managed file system on the file server, wherein data files migrated to the HSM server are recorded in a list of migrated data files having a unique identifier for each of the migrated data files, the method comprising the steps of:
- querying the list of migrated data files migrated from the managed file server to the HSM server;
 - for each file entry in the list of migrated data files, retrieving from the managed file system at least one attribute of the migrated data file that is identified by the corresponding unique identifier;
 - comparing the retrieved attributes with the corresponding attributes stored in the list of migrated data files; and
 - updating the HSM server for the migrated managed file system dependent on the results of the preceding step of comparing.
14. The method according to claim 13, wherein performing the steps of claim 13 by a reconciling process and wherein the reconciling process requests the list of migrated data files via the network from the HSM server.

15. A hierarchical storage management (HSM) system including at least one HSM server and at least one file server having stored a managed file system, the at least one HSM server and the at least one file server being interconnected via a network, where data files are migrated temporarily from the at least one file server to the at least one HSM, the system comprising:

means for prespecifying a scanning scope determined by a number of candidate data files

means for scanning the file system until having reached the prespecified number of migration candidate files and for identifying candidate data files to be migrated;

means for monitoring the managed file system;

means for migrating candidate data files to the HSM server;

means for reconciling the managed file system.

16. The system according to claim 15, further comprising a means for replacing a migrated data file in the managed file system by a stub file providing at least information about the location of the migrated data file on the HSM server.

17. The system according to claim 15, further comprising means for assigning a unique identifier to at least part of the candidate data files stored in the storage means.

18. The system according to claim 15, further comprising at least two storage means for identifying candidate data files, where the first storage means is generated and/or updated by a scanning process and where the at least second storage means is used by an automigration process, and where the automigration process gathers the content of the first storage means from the scanning process when all candidate data files of the at least second storage means are migrated.
19. A data processing program for execution in a data processing system comprising software code portions for performing a method comprising:
- providing at least one list for identifying candidate data files to be migrated;
 - prespecifying a scanning scope determined by the number of candidate data files;
 - scanning the managed file system until having reached the prespecified number of migration candidate data files;
 - selecting migration candidate data files according to at least one attribute;
 - recording the selected migration candidate data files in the provided at least one list for identifying candidate data files; and
 - migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server.

20. An article of manufacture comprising a program storage medium readable by a processor and embodying one or more instructions executable by the processor to perform a method comprising:
- providing at least one list for identifying candidate data files to be migrated;
 - prespecifying a scanning scope determined by a number of candidate data files;
 - scanning the managed file system until having reached the prespecified number of migration candidate data files;
 - selecting migration candidate data files according to at least one attribute;
 - recording the selected migration candidate data files in the provided at least one list for identifying candidate data files; and
 - migrating at least part of the selected candidate data files identified in the at least one list for identifying candidate data files from the file server to the HSM server.

9. Evidence Appendix

None

10. Related Proceedings Appendix

None